

AugRefer: Advancing 3D Visual Grounding via Cross-Modal Augmentation and Spatial Relation-based Referring

Xinyi Wang¹, Na Zhao^{2*}, Zhiyuan Han¹, Dan Guo³, Xun Yang¹

¹University of Science and Technology of China

²Singapore University of Technology and Design

³Hefei University of Technology

{wxy1, aaronhan,xyang21}@mail.ustc.edu.cn, na_zhao@sutd.edu.sg, guodan@hfut.edu.cn

Abstract

3D visual grounding (3DVG), which aims to correlate a natural language description with the target object within a 3D scene, is a significant yet challenging task. Despite recent advancements in this domain, existing approaches commonly encounter a shortage: a limited amount and diversity of text-3D pairs available for training. Moreover, they fall short in effectively leveraging different contextual clues (e.g., rich spatial relations within the 3D visual space) for grounding. To address these limitations, we propose AugRefer, a novel approach for advancing 3D visual grounding. AugRefer introduces cross-modal augmentation designed to extensively generate diverse text-3D pairs by placing objects into 3D scenes and creating accurate and semantically rich descriptions using foundation models. Notably, the resulting pairs can be utilized by any existing 3DVG methods for enriching their training data. Besides, AugRefer presents a language-spatial adaptive decoder that effectively adapts the potential referring objects based on the language description and various 3D spatial relations. Extensive experiments on three benchmark datasets clearly validate the effectiveness of AugRefer.

1 Introduction

3D visual grounding (3DVG) stands as an important and challenging task, aimed at locating objects within 3D scenes based on provided textual descriptions. As an advancement of 3D object detection (Zhao et al. 2020; 2022; Sheng et al. 2022; Han et al. 2024), it plays a critical perceptual role in various downstream applications, thus attracting increasing research attention. The emergence of large language models (LLMs) adds further allure to this field, offering a pathway to connect LLMs with the physical 3D world seamlessly.

Existing 3DVG methods commonly encounter a shortage of diverse training data pairs, consisting of the 3D scene with the referred object and the corresponding language description. This issue inherently arises from limitations on the 3D data side, where collecting and annotating 3D data is complex, costly, and time-consuming (Dai et al. 2017; Ding et al. 2023; Wang et al. 2023). For example, the popular 3DVG dataset (Achlioptas et al. 2020) only contains 1.5k scenes. Recent works (Hong et al. 2023; Zhang et al. 2023) have

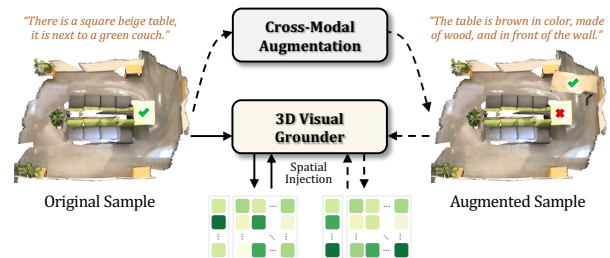


Figure 1: A brief illustration of our proposed AugRefer.

used LLMs to enrich linguistic descriptions but have not addressed 3D data scarcity, while other studies (Ge et al. 2024; Zhao et al. 2022; Zhang et al. 2020) have explored 3D augmentation to introduce objects into existing scenes. However, these single-modal augmentation techniques cannot be directly applied to cross-modal 3DVG due to two unique prerequisites inherent in augmenting text-3D pairs: 1) Ensuring **accurate correspondence** between 3D targets and linguistic descriptions, and 2) Providing **rich clues** necessary for locating the target, differentiating from other objects by incorporating both semantic and spatial information.

To tackle these limitations, we propose a novel approach for advancing 3D visual grounding, named AugRefer, as depicted in Fig. 1. In AugRefer, our initial step involves devising a cross-modal augmentation mechanism to enrich 3D scenes by injecting objects and furnishing them with diverse and precise descriptions. This augmentation process involves three main steps: inserting objects into 3D scenes, rendering these scenes into 2D images, and using foundation models to generate detailed captions. Furthermore, we design a multi-granularity rendering strategy to capture intricate textures and tailored prompts to produce diverse captions for each level. As a result, our cross-modal augmentation addresses the issue of data scarcity in 3DVG by significantly increasing the quantity and diversity of text-3D pairs.

In generated text-3D pairs, more complex situations arise. As shown in Fig. 1, if a scene already contains a table and our augmentation introduced an external table as the grounded target, the original one may become a distractor¹,

¹Distractors refers to objects of the same category as the target.

*Corresponding author.

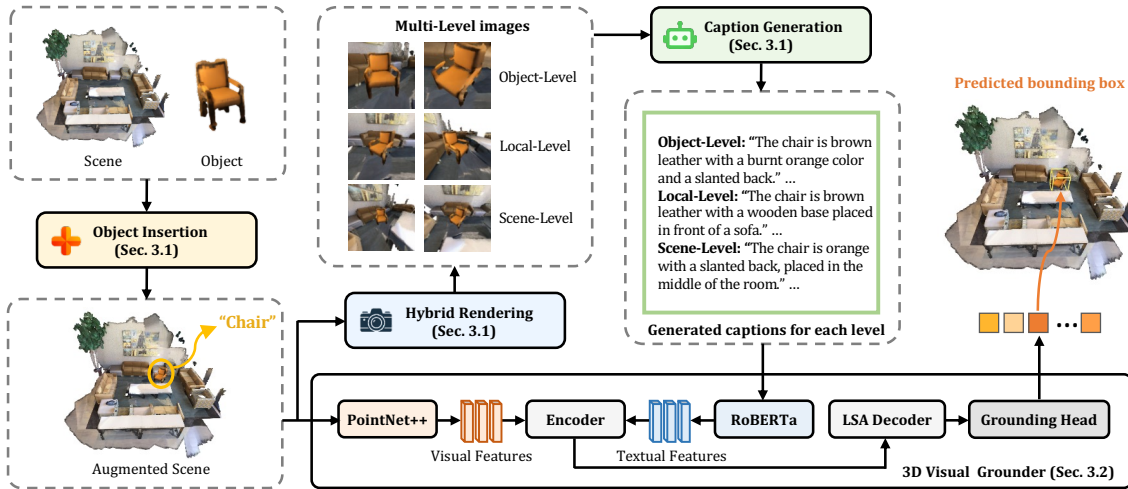


Figure 2: **The framework overview of AugRefer.** It consists of two components: 1) **Cross-Modal Augmentation** with three steps: ① Object Insertion → ② Hybrid Rendering → ③ Caption Generation; and 2) **3D Visual Grounder**, where our designed Language-Spatial Adaptive Decoder (LSAD) aims to enable more precise grounding by incorporating 3D spatial relations.

complicating the learning process. In such cases, it is essential to leverage spatial and other contextual information to make distinctions. To date, many 3DVG methods either overlook the exploitation of valuable contextual clues (Jain et al. 2022; Wu et al. 2023; Huang et al. 2022), such as rich spatial relations within the 3D visual space, or encounter challenges in effectively adapting the object features to different contextual clues (Zhao et al. 2021; Chen et al. 2022; Yang et al. 2024a). In our AugRefer, we design a Language-Spatial Adaptive Decoder (LSAD) as the cross-modal decoder to facilitate more accurate grounding of the target object. The LSAD is engineered to adapt the features of potential target objects (*i.e.*, the input to the decoder) to various contextual clues, including referring clues within the language description, object similarities in the 3D semantic space, and spatial relations within the 3D visual space. Our LSAD explores two distinct types of spatial relations within the 3D visual space: *global spatial relations* between objects and the entire scene as well as *pairwise spatial relations* between objects as illustrated in Fig. 5 (a). Furthermore, we inject these spatial relations into the attention mechanism within the decoder in a novel manner. It’s worth noting that our LSAD is compatible with any existing 3DVG framework employing a transformer-based architecture.

By integrating these two components, AugRefer achieves SOTA results on the ScanRefer and Sr3D datasets, showing the effectiveness of our proposed method. Furthermore, we demonstrate that our method can seamlessly integrate with existing 3DVG methods, such as BUTD-DETR and EDA, leading to consistent and significant improvements.

2 Related Work

3D Visual Grounding focuses on locating the language-referred object in 3D point clouds, which is different from the 2D visual grounding (Yang et al. 2021a, 2022). Ow-

ing to advances in transformers (Vaswani et al. 2017; Pan et al. 2024), the transformer-based methods have emerged as mainstream in 3DVG. Most methods use attention mechanisms to fuse multi-modal features implicitly. For example, BUTD-DETR (Jain et al. 2022) employs transformer-based encoder and decoder layers to fuse 3D visual features with features from other streams. EDA (Wu et al. 2023) performs more fine-grained alignment between visual and textual features by decoupling the input text. However, these standard attention modules lack the incorporation of spatial relationships. To address this issue, 3DVG-Transformer (Zhao et al. 2021) incorporates distance of proposals to capture pairwise spatial relations. CORE-3DVG (Yang et al. 2024a) exploits the spatial features under the guidance of linguistic cues. In this paper, we propose an effective spatial relation referring module for better global and pairwise perception.

3D Data Augmentation aims to mitigate the challenge of data scarcity and significantly enhance performance. Early efforts include but are not limited to, geometric transformations, noise injection, and generative methods. Several recent indoor augmentation techniques also follow the practice of mixing samples in 3D outdoor detection tasks. For instance, Mix3D (Nekrasov et al. 2021) directly merges two point cloud scenes to achieve scene-level data augmentation, while DODA (Ding et al. 2022) creatively implements a cuboid-level merge between source and target point clouds, tailored for domain adaptation scenarios. On the other hand, a few studies investigate the augmentation of 3D scenes with additional objects. For example, the outdoor 3D detection method Moca (Zhang et al. 2020) pastes ground-truth objects into both Bird’s Eye View (BEV) and image features of training frames. Likewise, 3D Copy-Paste (Ge et al. 2024) inserts virtual objects into real indoor scenes. In this work, we focus on implementing cross-modal augmentation between text descriptions and 3D scenes.




Figure 3: **a) Multi-Angle Camera:** For each level of the scene, images are captured from multiple angles. **b) Multi-Level Rendering:** The scene is rendered at different levels.

3 Methodology

Consider a 3D indoor scene denoted by a 3D point cloud P and an textual description T . Our goal is to predict the location $B_t \in \mathbb{R}^6$ of the targeted 3D object based on the description T . To tackle the scarcity of 3D object-text pairs and incorporate object-to-object as well as scene-wide spatial context into object grounding, we propose a novel method, AugRefer, that performs cross-modal augmentation and spatial relation-based referring, as shown in Fig. 2.

3.1 Cross-Modal Augmentation

Our goal is to significantly diversify the limited 3D object-text training pairs. Our proposed cross-modal augmentation is a plug-and-play solution, easily integrable into existing models. The process (as illustrated in Fig. 2) involves manipulating the 3D scenes in three key steps: **1) Insertion:** selecting suitable insertion positions for new objects and placing them in a plausible way that avoids interference with existing objects; **2) Rendering:** rendering the inserted objects in a multi-granularity way; and **3) Captioning:** taking the snapshots to generate diverse yet realistic descriptions. These captions are then refined to enhance their precision.

Object Insertion. For a 3D scene and an external object, where the latter is randomly selected from other scenes, our first consideration is *where*, *what*, and *how* to insert into the scene. To this end, we propose three main constraints for the Insertion operation: 1) *the ground plane*, 2) *the stander inserted object*, and 3) *no collision with existing objects in the 3D scene*. Algorithm 1 in the supplementary material outlines the plausible Insertion algorithm.

In this step, we designate the floor as the primary area for introducing new elements, selecting the insertion plane with the smallest Z-axis value. This entails imposing specific categorical constraints on the objects designated for insertion. Therefore, we focus on objects classified as a stander (*e.g.*, table and chair) which naturally stands on the ground plane, rather than a hanger (*e.g.*, window and curtain). Subsequently, following (Zhao et al. 2022), we simplify collision detection by converting the 3D scene into 2D floor images. Specifically, we employ an erosion technique in which

the size of the object’s shape determines the kernel used to erode the ground plane, thereby identifying the collision-free area suitable for insertion. If no suitable insertion area is found, we will resample another object and check the available space for insertion. This search process will continue until a viable insertion area is identified or the search limit is reached. Finally, before the actual insertion, the object undergoes random jittering, flipping, and rotation along the Z-axis. The most frequently inserted stander objects during this process are chairs, cabinets, and tables.

Hybrid Rendering. A crucial step in our cross-modal augmentation involves linking augmented 3D scenes with appropriate descriptions. We achieve this by projecting the 3D scenes into 2D images and then generating rich and accurate descriptions through image captioning. To ensure the generation of high-quality descriptions, precise and visually detailed 2D images are essential. Therefore, we employ a hybrid rendering strategy that considers both multi-angle and multi-view aspects, as illustrated in Fig. 3.

Firstly, we develop a multi-angle camera placement strategy to address the occlusions present in the 3D point cloud. These occlusions arise from the cluttered nature of scenes and the inherent limitations of point cloud data collection, which often lead to incomplete object capture. In our multi-angle camera placement strategy, we position cameras at 0, 45, and 90 degrees relative to the object’s center and rotate them around the object to obtain multiple snapshots. Secondly, we design a multi-level rendering strategy that encompasses object-, local-, and scene-level views to provide detailed attributes and spatial relationships of the inserted objects, as illustrated in Fig.3 (b). Specifically, we render images at three levels of detail by centering the inserted object and adjusting the field of view: 1) *object-level*: the object fills the frame, providing detailed insights into its categories and attributes. 2) *local-level*: with a broader view showing the object’s relationships with adjacent regions. 3) *scene-level*: the view is expanded to include almost the entire scene for global contextual information. Lastly, despite using multi-angle and multi-view rendering, issues like missing point clouds and obstructions can still arise and degrade image quality. To address this, we calculate CLIP (Radford et al. 2021) similarity scores between the images and their classes, selecting the top M images for the captioning phase.

Diverse Description Generation. Building upon the success of 2D multi-modal pre-trained models (Li et al. 2023; Achiam et al. 2023), we propose a strategy utilizing BLIP2 (Li et al. 2023) to generate accurate referential alignment. We meticulously craft various BLIP2 input prompts tailored for different levels of rendered images, as illustrated in Fig. 4. At the object-level, we instruct the model to provide detailed descriptions to capture finer visual characteristics. For the local and scene levels, we also require it to convey the spatial relationships between objects and their surroundings. Thus, we guide it to first identify the surrounding objects and then describe their interrelations. In order to refine the captions at each level, we employ GPT-3.5 (Brown et al. 2020) to automatically identify and rectify potential inaccuracies prior to summarization. Additionally, we use GPT-3.5

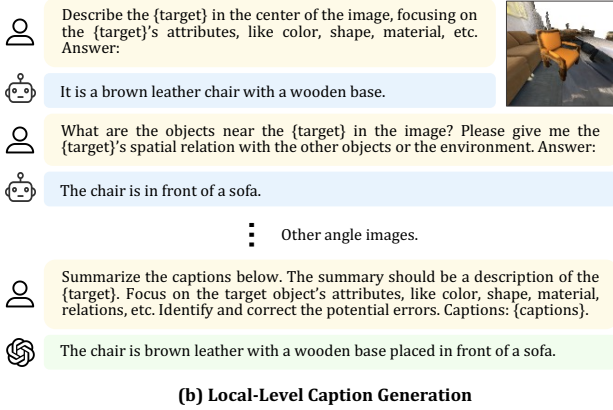
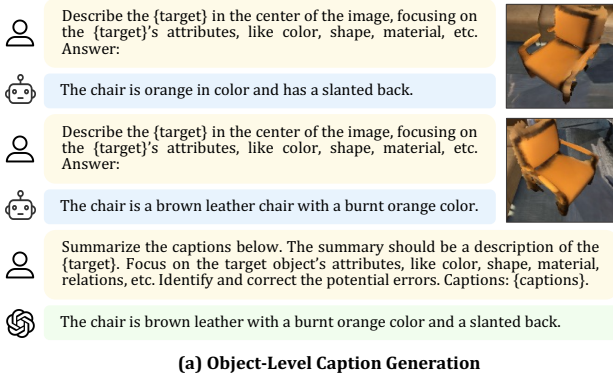


Figure 4: **Multi-Level Caption Generation.** Conversation process with BLIP2 and ChatGPT for captioning various level rendering images. Both the Local-Level and Scene-Level captions utilize the same set of prompts. We describe the approach using the Local-Level as an example.

to rephrase the captions, enhancing the diversity of descriptions and augmenting the textual modality.

3.2 Overview of 3D Visual Grounder

Following existing approaches (Jain et al. 2022; Wu et al. 2023), our 3D visual grounder consists of four basic modules: a feature extractor, a feature encoder, a cross-modal decoder, and a grounding head, as illustrated in Fig. 2.

Feature Extractor. We use a pre-trained PointNet++ (Qi et al. 2017) to encode the input point cloud and extract visual features $\mathcal{V} \in \mathbb{R}^{N_p \times d}$. We adopt a pre-trained RoBERTa (Liu et al. 2019) model to encode the textual input, generating language features $\mathcal{T} \in \mathbb{R}^{N_l \times d}$. Here, N_p and N_l denote the length of the visual and language tokens, respectively. Following (Jain et al. 2022), we use box stream extracted from the GroupFree (Liu et al. 2021) detector to provide bounding box guidance for the visual features. We then utilize a learnable MLP layer to transform the N_b detected bounding boxes into feature representation $\mathcal{B} \in \mathbb{R}^{N_b \times d}$.

Feature Encoder. Within the encoder, visual and language features interact through the standard cross-attention layers, where they cross-attend to each other and to box proposal

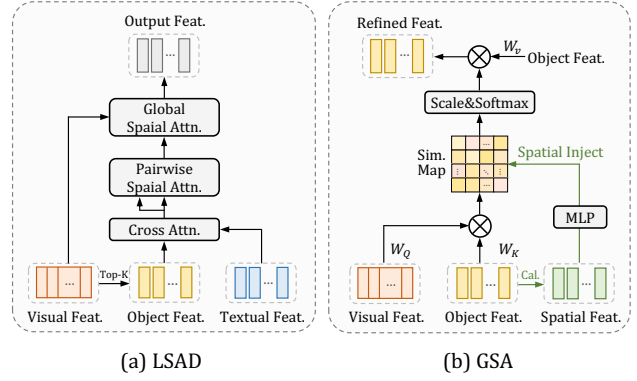


Figure 5: Illustrations of a) Language-Spatial Adaptive Decoder (LSAD) layer and b) Global Spatial Attention (GSA).

tokens using standard key-value attention in each layer. Following the acquisition of cross-modal features $F_v \in \mathbb{R}^{N_p \times d}$ and $F_t \in \mathbb{R}^{N_l \times d}$, a linear layer is employed to select the top K visual features, denoted as $F_o \in \mathbb{R}^{K \times d}$, representing target object candidates.

Cross-Modal Decoder. We design the decoder as a language-spatial adaptive decoder (LSAD), which further refines these candidates’ visual features with the guidance of various contextual clues from different sources such as language, box stream, and visual information. We will discuss the details of LSAD in Sec. 3.3.

Grounding Head. The output object features of the decoder are fed into an MLP layer, which predicts the referential object bounding box. Following our baselines (Jain et al. 2022; Wu et al. 2023), we respectively project visual and textual features into two linear layers, whose weights are denoted as $P_v \in \mathbb{R}^{K \times 64}$ and $P_t \in \mathbb{R}^{N_l \times 64}$, and compare the outputs with the ground-truths using soft-class prediction loss and semantic alignment loss.

3.3 Language-Spatial Adaptive Decoder

The cross-modal decoder within our 3D visual grounder stands as the most critical module, tasked with modeling the contextual relationships between the target object and relevant objects that align with the language description, thereby facilitating the identification of the correct referred object. Typically, the cross-modal decoder module, *e.g.*, the decoder in (Jain et al. 2022; Wu et al. 2023), utilizes cross-attention to capture the relationships between potential objects and language (or 3D object proposals) while applying self-attention to refine the features of candidate objects further. However, these conventional attention processes, which only model object-to-object relationships at a semantic level, cannot explicitly incorporate the spatial relationships between objects. Such spatial relationships are crucial for 3DVG (Zhao et al. 2021; Chen et al. 2022; Yang et al. 2024a), as the language description usually denotes objects based on their relative spatial positions within 3D scenes.

Furthermore, the inclusion of rich text-3D pairs generated by our cross-modal augmentation exacerbates the necessity of incorporating spatial relationships into the decoder. To

address this, we introduce a Language-Spatial Adaptive Decoder (LSAD) designed to incorporate spatial relations from both global and pairwise perspectives.

The architecture of our LSAD layer is illustrated in Fig. 5 (a). In each decoder layer, we employ three distinct types of attention to refine the visual features of the objects. We initially perform cross-attention between object features and textual features, assigning weights to objects based on their relevance to the language. Subsequently, the objects engage in pairwise spatial attention to aggregate relative spatial relationships, followed by global spatial attention to gather global cues. After N_D decoder layers, the final visual features are fed into the grounding head to predict the target.

Global Spatial Attention. We design this module to achieve a better understanding of scene-wide spatial context, considering that global position descriptions also appear in the dataset, especially in our global-level augmented annotations, such as “*This speaker is brown in the corner.*” and “*This nightstand is in the middle.*”. Therefore, we introduce global spatial attention, which injects spatial relation information in the same manner as in pairwise spatial attention. However, the calculation of global spatial relationships and the attention targets differ. Specifically, we calculate the normalized coordinates of the object center in the entire scene as global spatial features $R_g \in \mathbb{R}^{K \times 1 \times d_g}$:

$$r_i^g = [x_{\text{norm}}, y_{\text{norm}}, z_{\text{norm}}]. \quad (1)$$

Then we transform the global spatial relationship R_g as F_g :

$$F_g = \text{MLP}(R_g). \quad (2)$$

To improve the integration of global spatial relations, we refine object characteristic features F_o by incorporating spatial features F_g and the visual features of the entire scene point cloud F_v . The process of global spatial attention (GSA) is outlined as follows:

$$Q = F_o W_Q, K = F_v W_K, V = F_v W_V, S_g = F_g W_S, \quad (3)$$

$$\text{GSA}(Q, K, V) = \text{softmax} \left(\frac{QK^T + S_g}{\sqrt{2d_h}} \right) V, \quad (4)$$

where W_Q, W_K, W_V, W_S denote learnable linear layers.

Pairwise Spatial Attention. In natural language descriptions, it is often necessary to distinguish the target from distractors by referring to one or more anchors and their pairwise spatial relationships, such as “*The chair next to a brown couch*” or “*There is a wooden cabinet between a water cooler and a trash can.*”. Therefore, we introduce Pairwise Spatial Attention (PSA), injecting spatial features into the same method as Global Spatial Attention. The calculation method of global relations and the features involved in the attention differ. Specifically, we calculate the distances and directions between objects to obtain the pairwise spatial relationships $R_p \in \mathbb{R}^{K \times K \times d_p}$, for K objects, where r_{ij}^p denotes the spatial relation between objects O_i and O_j (see the supplementary material for more details).

4 Experiments

4.1 Dataset and Experimental Setting

Datasets. We use three 3DVG datasets: ScanRefer (Chen et al. 2020), Nr3D (Achlioptas et al. 2020), and Sr3D (Achlioptas et al. 2020) to evaluate our method. Note that Sr3D and

Nr3D provide ground-truth objects, and some methods simplify the grounding task to a matching problem of selecting the ground-truth box that best matches the description. Following (Yang et al. 2024a), we use detected objects as input with the raw point cloud, instead of ground truths.

Evaluation Metrics. We evaluate performance using the Acc@k (k=0.25 or 0.5) metric. Acc@k means the accuracy where the best-matched proposal has an intersection over union with the ground truth greater than the threshold k.

Baselines. We choose two 3DVG models, specifically BUTD-DETR (Jain et al. 2022) and EDA (Wu et al. 2023) as baselines². In our experiments, we integrate our cross-modal augmentation and hierarchy spatial decoder into the baselines, respectively, to verify the effectiveness of our method.

Implementation Details. Our experiments are conducted on four NVIDIA A100 80G GPUs, utilizing PyTorch and the AdamW optimizer. We adjust the batch size to 12 or 48 and augment training with 22.5k generated pairs for each dataset. On average, the generated description contains 13.7 words. The visual encoder’s learning rate is set to 2e-3 for ScanRefer, while other layers are set to 2e-4 across 150 epochs. In contrast, SR3D and NR3D have learning rates of 1e-3 and 1e-4, respectively; NR3D undergoes 200 epochs of training, whereas SR3D requires only 100 epochs due to its simpler, template-generated descriptions.

4.2 Overall Comparison

In Tab. 1 and Tab. 2, we report the comparison of performance between the baselines and our AugRefer, as well as the comparison between ours and the reported SOTA results, across three 3DVG datasets, *i.e.* ScanRefer, Nr3D, and Sr3D. This leads us to the following insights:

- Our method exhibit significant improvements across various metrics when integrated into the baselines, BUTD-DETR and EDA. In particular, when compared to BUTD-DETR, our performance improvements *w.r.t.* the overall Acc@0.25 metric are 3.05%, 9.81%, and 6.58%, respectively on ScanRefer, Nr3D, and Sr3D. Furthermore, by integrating EDA, the open-source model with the best performance, AugRefer significantly increases accuracy by 2.10%, 4.41%, and 6.56% across the three datasets.
- Benefiting from our AugRefer, EDA has achieved the SOTA level in Nr3D and Sr3D datasets. In Scanrefer, we have brought EDA closer to the level of SOTA. It’s noteworthy that our AugRefer is also compatible with the SOTA model CORE-3DVG (Yang et al. 2024a), offering the potential for further performance enhancements.

4.3 In-depth Studies

Two strategies synergize to make AugRefer effective. The results in Tab. 3 show that our cross-model augmentation significantly boosts performance in the simple splits, *i.e.*, ‘*unique*’, owing to the inclusion of augmented 3D scenes and a broader array of augmented objects, which enhance

²Since the SOTA method CORE-3DVG is not open source, we use the top-performing models EDA and BUTD-DETR in 3DVG. Note that our method could be compatible with CORE-3DVG.

Method	unique		multiple		overall	
	Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5
ReferIt3D (Achlioptas et al. 2020)	53.75	37.47	21.03	12.83	26.44	16.90
ScanRefer (Chen, Chang, and Nießner 2020)	67.64	46.19	32.06	21.26	38.97	26.10
TGNN (Huang et al. 2021)	68.61	56.80	29.84	23.18	37.37	29.70
SAT (Yang et al. 2021b)	73.21	50.83	37.64	25.16	44.54	30.14
FFL-3DOG (Feng et al. 2021)	78.80	67.94	35.19	25.70	41.33	34.01
InstanceRefer (Yuan et al. 2021)	77.45	66.83	31.27	24.77	40.23	32.93
3DVG-Transformer (Zhao et al. 2021)	77.16	58.47	38.38	28.70	45.90	34.47
Multi-View Transformer (Huang et al. 2022)	77.67	66.45	31.92	25.26	40.80	33.26
3D-SPS (Luo et al. 2022)	81.63	64.77	39.48	29.61	47.65	36.43
Vi13DRef (Chen et al. 2022)	81.58	68.62	40.30	30.71	47.94	37.73
ViewRefer (Guo et al. 2023)	76.35	64.27	33.08	26.50	41.35	33.69
CORE-3DVG (Yang et al. 2024a)	84.99	67.09	<u>51.82</u>	<u>39.76</u>	<u>56.77</u>	43.83
BUTD-DETR (Jain et al. 2022)	81.55	64.39	44.81	33.41	50.86	38.51
+AugRefer	85.21 +3.66	68.99 +4.60	47.73 +2.92	37.16 +3.75	53.91 +3.05	42.41 +3.90
EDA (Wu et al. 2023)	83.81	64.62	47.91	36.29	53.58	40.77
+AugRefer	86.21 +2.40	70.75 +6.13	49.96 +2.05	39.06 +2.77	55.68 +2.10	44.03 +3.26

Table 1: Comparison with SOTA methods on *ScanRefer*. We highlight the best performance with underlining.

Method	Nr3D	Sr3D
	Acc@0.25	Acc@0.25
ReferIt3D (Achlioptas et al. 2020) †	24.00	27.70
InstanceRefer (Yuan et al. 2021) †	29.90	31.50
LanguageRefer (Roh et al. 2022) †	28.60	39.50
SAT (Roh et al. 2022) †	31.70	35.40
CORE-3DVG (Yang et al. 2024a)	<u>49.57</u>	54.30
BUTD-DETR (Jain et al. 2022)	38.60	53.64
+AugRefer	48.41 +9.81	60.22 +6.58
EDA (Wu et al. 2023)	42.08	51.39
+AugRefer	46.49 +4.41	57.95 +6.56

Table 2: Comparison with SOTA methods on *Nr3D* and *Sr3D*. † Evaluation results are quoted from (Yang et al. 2024a). We highlight the best performance with underlining.

	Method		unique		multiple		overall	
	CA	LSAD	@0.25	@0.5	@0.25	@0.5	@0.25	@0.5
(a)			81.55	64.39	44.81	33.41	50.86	38.51
(b)	✓		85.13	66.62	46.40	34.96	52.79	40.18
(c)		✓	84.52	67.92	46.85	36.03	53.07	41.29
(c)	✓	✓	85.21	68.99	47.73	37.16	53.91	42.41

Table 3: The ablation study of our AugRefer. CA stands for cross-modal augmentation; LSAD for language-spatial adaptive decoder. BUTD-DETR is used as baseline (row a).

the model’s ability to perceive object classes. In contrast, spatial injection yields greater improvements in the more challenging splits, *i.e.*, ‘*multiple*’. This stems from the learning of the injected spatial relations, thereby strengthening the model’s capability to differentiate distractors according to spatial relations. Together, these two strategies synergize to form the rational AugRefer. Overall, these findings highlight the complementary benefits of our proposed ap-

	object	Level		overall	
		local	scene	@0.25	@0.5
(a)				50.86	38.51
(b)	✓			48.34	36.67
(c)		✓		49.32	37.06
(d)			✓	50.38	38.74
(e)		✓	✓	51.04	38.70
(f)	✓	✓	✓	52.79	40.18

Table 4: The ablation study of augmentation levels.

	Quantity per scene and level	overall	
		@0.25	@0.5
(a)	1	51.49	39.15
(b)	3	52.79	40.18
(c)	5	51.37	38.38

Table 5: The ablation study on augmentation quantities.

proaches in improving 3D visual grounding performance. **Multi-level enhances accuracy.** We apply multi-level rendering and caption generation strategies in our cross-modal augmentation. To investigate the effect of different levels in the cross-augmentation, we gradually introduce distinct levels of augmented samples into the baseline BUTD-DETR and report the results in Tab. 4. It is interesting to note that while each level of augmentation alone does not improve overall performance compared to the baseline without augmentation (row a), combining all three levels of granularity in augmentation allows AugRefer to achieve improvements of 1.93% and 1.67% in overall metrics. This highlights the importance of our multi-level design, which captures detailed attributes and spatial relationships of newly inserted objects, providing more precise descriptions. In addition, we also conduct experiments to investigate the rela-




Figure 6: Qualitative results with ScanRefer descriptions

	Method	overall	
		@0.25	@0.5
(a)	pair-glb-lang	51.18	39.27
(b)	pair-lang-glb	52.93	41.02
(c)	lang-pair-glb	53.91	42.41

Table 6: The ablation study on language priors and visual priors in the LSAD module.

relationship between the quantity of generated pairs and performance improvements. During the training phase of BUTD-DETR baseline, we randomly add n pairs ($n = 1, 3, 5$) from each scene and level to augment the training dataset. We observed that increasing the number of generated pairs can lead to performance degradation, likely due to noise introduced in the generated pairs. To maintain a balance between generated and original pairs, we set n to 3.

Our current LSAD design outperforms the alternatives.

In Tab. 6, we investigate the effects of order in applying three types of attention on ScanRefer dataset with cross-modal augmented training data. Rows (a-c) incorporate the LSAD module with varying orders of the three attentions. The results in Tab. 6 show that modeling complex spatial relationships is more effective when language priors are applied (cross-attended) first. Besides, we compare our LSAD module with an alternative design that also models spatial relationships to highlight the advantages of our LSAD. Specifically, we replace the LSAD decoder in our baseline BUTD-DETR with the corresponding structure from Vi3DRef (Chen et al. 2022), denoted as ‘ViL. Decoder’ in Tab. 7. The comparison results show that our LSAD module, which integrates both global and pairwise relationships, achieves significantly better performance.

	Method	overall	
		@0.25	@0.5
(a)	BUTD-DETR	50.86	38.51
(b)	+ViL. Decoder	50.18	38.65
(c)	+LSAD	53.91	42.41

Table 7: Comparison of the spatial relation module.

Qualitative Analysis: Fig. 6 illustrates the qualitative comparison between AugRefer and the baseline BUTD-DETR using two samples from ScanRefer dataset. The visualization results demonstrate that our AugRefer outperforms the baseline, particularly in grounding challenging objects such as the small box. Furthermore, AugRefer exhibits notable improvements in spatial modeling, particularly in the pairwise (first three columns) and global (last column) contexts.

5 Conclusion

Our work alleviates the shortage of both amount and diversity in the text-3D grounding dataset and the inefficiencies of exploring contextual clues by introducing AugRefer. We enrich 3D scenes with additional objects and generate detailed descriptions in three distinct granularities using foundation models. Moreover, we integrate contextual clues into the model, enabling a thorough comprehension of these relationships. This paper pioneers the use of cross-modal augmentation techniques, substantially advancing the field of 3D visual grounding and providing viable solutions wherever in research and practical fields. In the future, we will try to extend our idea to tackle more complex reasoning tasks (Yang et al. 2024b; Luo et al. 2025).

Acknowledgments

This research work was supported by the National Natural Science Foundation of China (NSFC) under Grant U22A2094, and also supported by the Agency for Science, Technology and Research (A*STAR) under its MTC Programmatic Funds (Grant No. M23L7b0021). We also acknowledge the support of the advanced computing resources provided by the Supercomputing Center of the USTC, and the support of GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Achlioptas, P.; Abdelreheem, A.; Xia, F.; Elhoseiny, M.; and Guibas, L. 2020. Referit3d: Neural listeners for fine-grained 3d object identification in real-world scenes. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, 422–440. Springer.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*.
- Chen, D. Z.; Chang, A. X.; and Nießner, M. 2020. Scanrefer: 3d object localization in rgb-d scans using natural language. In *European conference on computer vision*, 202–221. Springer.
- Chen, S.; Guhur, P.-L.; Tapaswi, M.; Schmid, C.; and Laptev, I. 2022. Language conditioned spatial relation reasoning for 3d object grounding. *Advances in neural information processing systems*, 35: 20522–20535.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5828–5839.
- Ding, R.; Yang, J.; Jiang, L.; and Qi, X. 2022. Doda: Data-oriented sim-to-real domain adaptation for 3d semantic segmentation. In *European Conference on Computer Vision*, 284–303. Springer.
- Ding, R.; Yang, J.; Xue, C.; Zhang, W.; Bai, S.; and Qi, X. 2023. Pla: Language-driven open-vocabulary 3d scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7010–7019.
- Feng, M.; Li, Z.; Li, Q.; Zhang, L.; Zhang, X.; Zhu, G.; Zhang, H.; Wang, Y.; and Mian, A. 2021. Free-form description guided 3d visual graph network for object grounding in point cloud. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3722–3731.
- Ge, Y.; Yu, H.-X.; Zhao, C.; Guo, Y.; Huang, X.; Ren, L.; Itti, L.; and Wu, J. 2024. 3D Copy-Paste: Physically Plausible Object Insertion for Monocular 3D Detection. *Advances in Neural Information Processing Systems*, 36.
- Guo, Z.; Tang, Y.; Zhang, R.; Wang, D.; Wang, Z.; Zhao, B.; and Li, X. 2023. Viewrefer: Grasp the multi-view knowledge for 3d visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Han, Y.; Zhao, N.; Chen, W.; Ma, K. T.; and Zhang, H. 2024. Dual-Perspective Knowledge Enrichment for Semi-supervised 3D Object Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Hong, Y.; Zhen, H.; Chen, P.; Zheng, S.; Du, Y.; Chen, Z.; and Gan, C. 2023. 3d-llm: Injecting the 3d world into large language models. *Advances in Neural Information Processing Systems*, 36: 20482–20494.
- Huang, P.-H.; Lee, H.-H.; Chen, H.-T.; and Liu, T.-L. 2021. Text-guided graph neural networks for referring 3d instance segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 1610–1618.
- Huang, S.; Chen, Y.; Jia, J.; and Wang, L. 2022. Multi-view transformer for 3d visual grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15524–15533.
- Jain, A.; Gkanatsios, N.; Mediratta, I.; and Fragkiadaki, K. 2022. Bottom up top down detection transformers for language grounding in images and point clouds. In *European Conference on Computer Vision*, 417–433. Springer.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Liu, Z.; Zhang, Z.; Cao, Y.; Hu, H.; and Tong, X. 2021. Group-free 3d object detection via transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2949–2958.
- Luo, C.; Di, D.; Yang, X.; Ma, Y.; Xue, Z.; Wei, C.; and Liu, Y. 2025. TrAME: Trajectory-Anchored Multi-View Editing for Text-Guided 3D Gaussian Splatting Manipulation. *IEEE Transactions on Multimedia*.
- Luo, J.; Fu, J.; Kong, X.; Gao, C.; Ren, H.; Shen, H.; Xia, H.; and Liu, S. 2022. 3d-sps: Single-stage 3d visual grounding via referred point progressive selection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16454–16463.
- Nekrasov, A.; Schult, J.; Litany, O.; Leibe, B.; and Engelmann, F. 2021. Mix3d: Out-of-context data augmentation for 3d scenes. In *2021 international conference on 3d vision (3dv)*, 116–125. IEEE.
- Pan, H.; Cao, Y.; Wang, X.; Yang, X.; and Wang, M. 2024. Finding and Editing Multi-Modal Neurons in Pre-Trained Transformers. In *Findings of the Association for Computational Linguistics ACL 2024*, 1012–1037.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30.

- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Roh, J.; Desingh, K.; Farhadi, A.; and Fox, D. 2022. Language-refer: Spatial-language model for 3d visual grounding. In *Conference on Robot Learning*, 1046–1056. PMLR.
- Sheng, H.; Cai, S.; Zhao, N.; Deng, B.; Huang, J.; Hua, X.-S.; Zhao, M.-J.; and Lee, G. H. 2022. Rethinking IoU-based optimization for single-stage 3D object detection. In *European Conference on Computer Vision*, 544–561. Springer.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, Z.; Huang, H.; Zhao, Y.; Li, L.; Cheng, X.; Zhu, Y.; Yin, A.; and Zhao, Z. 2023. Distilling coarse-to-fine semantic matching knowledge for weakly supervised 3d visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2662–2671.
- Wu, Y.; Cheng, X.; Zhang, R.; Cheng, Z.; and Zhang, J. 2023. Eda: Explicit text-decoupling and dense alignment for 3d visual grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19231–19242.
- Yang, L.; Zhang, Z.; Qi, Z.; Xu, Y.; Liu, W.; Shan, Y.; Li, B.; Yang, W.; Li, P.; Wang, Y.; et al. 2024a. Exploiting Contextual Objects and Relations for 3D Visual Grounding. *Advances in Neural Information Processing Systems*, 36.
- Yang, X.; Feng, F.; Ji, W.; Wang, M.; and Chua, T.-S. 2021a. Deconfounded video moment retrieval with causal intervention. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1–10.
- Yang, X.; Wang, S.; Dong, J.; Dong, J.; Wang, M.; and Chua, T.-S. 2022. Video moment retrieval with cross-modal neural architecture search. *IEEE Transactions on Image Processing*, 31: 1204–1216.
- Yang, X.; Zeng, J.; Guo, D.; Wang, S.; Dong, J.; and Wang, M. 2024b. Robust Video Question Answering via Contrastive Cross-Modality Representation Learning. *SCIENCE CHINA Information Sciences*, 67: 1–16.
- Yang, Z.; Zhang, S.; Wang, L.; and Luo, J. 2021b. Sat: 2d semantics assisted training for 3d visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1856–1866.
- Yuan, Z.; Yan, X.; Liao, Y.; Zhang, R.; Wang, S.; Li, Z.; and Cui, S. 2021. Instancerefer: Cooperative holistic understanding for visual grounding on point clouds through instance multi-level contextual referring. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1791–1800.
- Zhang, W.; Wang, Z.; and Loy, C. C. 2020. Exploring data augmentation for multi-modality 3d object detection. *arXiv preprint arXiv:2012.12741*.
- Zhang, Y.; Gong, Z.; and Chang, A. X. 2023. Multi3drefer: Grounding text description to multiple 3d objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15225–15236.
- Zhao, L.; Cai, D.; Sheng, L.; and Xu, D. 2021. 3DVG-Transformer: Relation modeling for visual grounding on point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2928–2937.
- Zhao, N.; Chua, T.-S.; and Lee, G. H. 2020. Sess: Self-ensembling semi-supervised 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11079–11087.
- Zhao, N.; and Lee, G. H. 2022. Static-dynamic co-teaching for class-incremental 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 3436–3445.
- Zhao, Y.; Zhao, N.; and Lee, G. H. 2022. Synthetic-to-Real Domain Generalized Semantic Segmentation for 3D Indoor Point Clouds. *arXiv preprint arXiv:2212.04668*.